

# From Conversation to Automation: Leveraging LLMs for Problem-Solving Therapy Analysis

Elham Aghakhani<sup>1</sup>, Lu Wang<sup>2</sup>, Karla T. Washington<sup>3</sup>  
George Demiris<sup>4</sup>, Jina Huh-Yoo<sup>2</sup>, Rezvaneh Rezapour<sup>1</sup>

<sup>1</sup>Drexel University, {ea664, sr3563}@drexel.edu

<sup>2</sup>Stevens Institute of Technology, {lwang97, jhuhyoo}@stevens.edu

<sup>3</sup>Washington University, kwashingt@wustl.edu

<sup>4</sup>University of Pennsylvania, gdemiris@nursing.upenn.edu

## Abstract

Problem-Solving Therapy (PST) is a structured psychological approach that helps individuals manage stress and resolve personal issues by guiding them through problem identification, solution brainstorming, decision-making, and outcome evaluation. As mental health care increasingly adopts technologies like chatbots and large language models (LLMs), it is important to thoroughly understand how each session of PST is conducted before attempting to automate it. We developed a comprehensive framework for PST annotation using established PST Core Strategies and a set of novel Facilitative Strategies to analyze a corpus of real-world therapy transcripts to determine *which* strategies are most prevalent. Using various LLMs and transformer-based models, we found that GPT-4o outperformed all models, achieving the highest accuracy (0.76) in identifying all strategies. To gain deeper insights, we examined *how* strategies are applied by analyzing Therapeutic Dynamics (autonomy, self-disclosure, and metaphor), and linguistic patterns within our labeled data. Our research highlights LLMs' potential to automate therapy dialogue analysis, offering a scalable tool for mental health interventions. Our framework enhances PST by improving accessibility, effectiveness, and personalized support for therapists.

## 1 Introduction

In recent years, modern stressors like COVID-19, global conflicts, and climate disasters have led to rising mental health conditions, especially among adults 35-44 (APA, 2023). In response, the World Health Organization (WHO) calls for urgent transformation in mental health care (WHO, 2022). Problem-Solving Therapy (PST) helps manage these stressors by enhancing problem-solving skills, reducing psychological distress, and strengthening coping mechanisms (Song and Sun, 2019; Arian et al., 2008; Gellis and Kenaley, 2008).

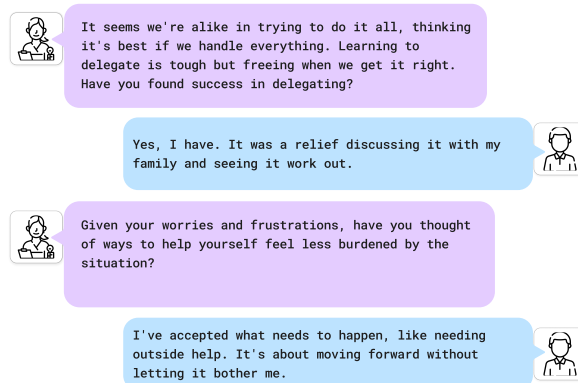


Figure 1: A real-life example of therapist and client exchanges in a problem-solving therapy session.

As rapidly evolving trends continue to shape the critical area of mental health, systematic methods for analyzing PST interactions remain limited (Chiu et al., 2024). This gap is especially relevant with the rise of large language model (LLM)-powered mental health chatbots (Lawrence et al., 2024; Cho et al., 2023). Before widespread adoption, it is crucial to assess how PST strategies are applied in practice, as therapists' approaches vary based on client needs and contexts (Andrade-Arenas et al., 2024). PST follows a structured process, identifying problems, brainstorming solutions, decision-making, and evaluation, to strengthen clients' problem-solving skills and coping abilities (Malouff et al., 2007). While new technologies hold promise, their adaptability and responsiveness must be ensured for effective PST automation (Izadi and Forouzanfar, 2024).

Previous work leveraged LLMs for analyzing therapist-client interactions (Lin et al., 2025; Demszky et al., 2023). For instance, Lee et al. used LLMs, particularly ChatGPT, to identify utterance-level features in therapist-client conversations to classify counselor's utterances. In PST, such research in automated annotations is still lacking due to limited data and costly domain expert annotation. Automating the annotation of various cogni-

tive therapy techniques is crucial for developing efficient models, like LLMs, that can navigate the complexity of therapeutic dialogues and identify effective interventions. This capability could enable real-time suggestions to therapists or direct interventions in automated systems.

Our work addresses this gap by using LLMs to automate PST annotation, using both closed- and open-source models to categorize therapist utterances. We analyzed 240 real-world anonymized PST sessions (68,306 dialogue exchanges), applying LLMs to annotate strategies based on existing and adapted PST and Facilitative strategy codebooks. Our results show that LLMs efficiently and accurately annotate PST strategies. We also analyzed how these strategies evolve across sessions finding that therapists progressively shift from exploratory strategies, such as defining problems and brainstorming solutions, to strategies focused on implementation and evaluation. Additionally, we analyze therapeutic dynamics (autonomy, self-disclosure, metaphor) and linguistic properties in all utterances to understand *how* these strategies were applied. Our results show that therapists adapt their language to session goals, using more non-directive language and metaphors in the early stages and transitioning to directive language and concrete communication as therapy progresses.

This approach enhances therapist-client interaction analysis, supports scalable and adaptive mental health interventions, and enables real-time identification of effective therapeutic techniques. Our framework advances NLP and mental health care, showing LLMs' potential to transform PST research. Finally, integrating Facilitative strategies refines the PST framework, offering deeper insights into targeted therapeutic interventions. All codes used in this study are available in our GitHub repository.<sup>1</sup>

## 2 Related Work

### 2.1 LLMs for Mental Health Care

LLMs have been applied for mental health care, leveraging distinct features such as conversational capabilities (Ma et al., 2023; Zheng et al., 2023), content generation (Yang et al., 2024; Smith et al., 2023), and detection of information patterns (Posner et al., 2011; Mauriello et al., 2021). Specifically, LLMs have been applied by conversational agents (Bendig et al., 2022) to serve as personal

digital companions (Ma et al., 2023), provide emotional support (Zheng et al., 2023), and support online counseling (Sharma et al., 2023; Filienko et al., 2024). Researchers applied LLMs to generate explanations for mental health analysis (Yang et al., 2024), help people learn mindfulness (Kumar et al., 2023), and create hypothetical case vignettes to social psychiatry (Smith et al., 2023). LLMs have also been used in the classification of subtypes of suicide (Posner et al., 2011), and the identification of sources of stress span (Mauriello et al., 2021).

However, gaps remain in previous research in accessing data with high external validity to fully understand the potential LLMs for different mental health care settings (Hua et al., 2024; Wang et al., 2021a; Goel et al., 2023). For example, among 36 datasets identified in Hua et al. (2024), 55.6% were collected from social media, 13.9% were simulated by clinicians, crowd workers, and LLMs, and only 8.3% came from real-life counseling conversations. Moreover, labels in existing datasets are typically not validated by experts (Hua et al., 2024; Harrigian et al., 2020). The absence of relevant evaluation frameworks brings further challenges. Every cognitive and interpersonal therapy technique has a unique, specialized structure and requirement for evaluating fidelity of delivery (Sharma et al., 2023; Mehta et al., 2022). Examining how professionals deliver interventions through therapy techniques is essential for addressing these gaps.

### 2.2 LLMs as (Co-)Annotators

Acquiring high-quality labeled data is time-consuming and costly when relying solely on human experts for annotation (Wang et al., 2021b, 2024; Bouzoubaa et al., 2024b). LLMs have been explored to support annotations through prompt engineering (Reynolds and McDonell, 2021; Goel et al., 2023; Bouzoubaa et al., 2024a), such as zero-shot, without any in-domain training (Liu et al., 2024; Reynolds and McDonell, 2021) and few-shot learning, incorporating examples of the task (Brown et al., 2020). Research showed LLMs outperformed crowd workers on annotation tasks including relevance, stance, topics, frame detection, and categorizing sentence segments (Gilardi et al., 2023; He et al., 2024; Törnberg, 2023; Pesianzadeh and Rezapour, 2025).

On the other hand, Zhu et al. (2023) and Ziemis et al. (2024) showed that the performance of LLMs varies depending on tasks, datasets, and labels, claiming the irreplaceableness of human annota-

<sup>1</sup>[github.com/social-nlp-lab/LLM\\_Problem\\_Solving\\_Therapy](https://github.com/social-nlp-lab/LLM_Problem_Solving_Therapy)

tors. Wang et al. (2021b) and Li (2024) showed that human-LLM collaboration, either through human annotation of LLM-labeled data or combining both annotations, outperformed LLM-only or human-only approaches. Li et al. (2023b) explored the allocation between humans and LLMs for annotation and contributed to a paradigm using uncertainty to estimate LLMs' annotation capability. Researchers also developed human-LLM collaborative annotation systems to support the annotations (Kim et al., 2024; Tang et al., 2024). Our study expands on previous research by employing LLMs to annotate therapy transcripts, based on experts' input, thereby refining existing frameworks and models and informing broader applications of LLMs in mental health care research.

### 3 Method

#### 3.1 Data

**Data Collection.** The PISCES trial (Demiris et al., 2019) provided transcriptions of 240 PST sessions between 152 family caregivers of older adults in hospice care and a behavioral interventionist. Each caregiver attended up to three therapy sessions, but not every session was recorded and transcribed due to technical and legal constraints. We first ensured the anonymity of these transcripts by removing any remaining information that could potentially reveal the identity of the clients or therapists. We then converted the transcripts into 68,306 dialogue exchanges between clients and therapists. An example of two exchanges is shown in Figure 1. We then separated the utterances<sup>2</sup> made by therapists, which resulted in 34,273 utterances. After consulting with domain experts and to ensure high quality and relevance, we excluded utterances shorter than 5 words, such as “hmmm”, “okay”, “well”, and “oh good” assuming that longer ones would provide more meaningful content for our study. This process resulted in 14,417 utterances for further analysis. The average word count for these utterances was 29.78 with a standard deviation of 42.72.

**Problem-Solving Therapy Codebook.** In collaboration with PST experts, we developed the annotation codebook using the core strategies from the ADAPT model of PST, which includes five steps: 1) developing a positive Attitude, 2) Defining the problem and setting goals, 3) being creative and

<sup>2</sup>An utterance is a continuous piece of speech, by one person, before or after which there is silence on the part of the person: <https://en.wikipedia.org/wiki/Utterance>

finding Alternative solutions, 4) Predicting how the solutions will work and making a plan, and 5) Trying out the solution plan (Demiris et al., 2019). We refer to these as the ‘Problem-Solving (PS) Core’ category. Upon further exploration of our data and various frameworks and codebooks for therapy conversations and strategies (Li et al., 2023a; Ribeiro et al., 2013; Wang et al., 2023; Ghamari Kivi and Jamshiddoust Miyanroudi, 2022), our domain experts suggested adding a new category of four strategies to this list, called ‘Facilitative’ strategies, which we will use interchangeably as ‘Facilitators’ throughout the paper. These strategies are further explained in Table 6.

**Data Annotation.** To develop an evaluation set, we randomly selected 500 therapist utterances for detailed annotation. One domain expert alongside one researcher from our team performed the annotation. The annotators were instructed to select at most one strategy from the ‘PS Core’ category and one strategy from the ‘Facilitators’, if applicable. If none of the strategies were suitable, we labeled it as ‘None.’ We used Cohen’s Kappa to assess annotators’ agreement across different classes of PST strategies. Our findings show that the agreement scores ranged from 0.69 to 0.88 for strategies in our codebook, indicating a substantial agreement among the annotators, reflecting a reasonable level of consistency in identifying the strategies. Agreement was higher for more straightforward classes (e.g., ‘Generating Alternative Solutions’) and lower for more subjective or nuanced classes (e.g., ‘Therapeutic Engagement’).

#### 3.2 PST Classification with LLM Prompting

LLMs, such as GPT-4, have been used for data annotation, offering capabilities for analyzing and categorizing diverse datasets (Li et al., 2023a; Kuzman et al., 2023; Hoes et al., 2023). In this work, we employed these models to annotate therapist utterances, aiming to identify the PST strategy used in each case. Specifically, we used two versions of OpenAI’s models: gpt-4o-2024-08-06 (OpenAI) (referred to as GPT-4o) and gpt-4-0125-preview (OpenAI, 2024) (referred to as GPT-4), as well as two open-source models, Llama-3.1-70B-Instruct (Meta) (referred to as Llama) and Yi-1.5-34B (Young et al., 2024) (referred to as Yi-1.5). All models were deployed with a zero-temperature setting to ensure deterministic output, enhancing the consistency of our annotations (See Appendix B for hy-

perparameters). Using this approach, we compared the efficiency and accuracy of these models against our 500 annotated utterances. We used two approaches to test how well LLMs can annotate data:

- **No Context:** We provided specific instructions to LLMs, followed by definitions of each strategy and a few examples in the prompts. We then presented each therapist’s utterance for classification, without providing any additional conversational context.
- **Two Previous Utterances as Context:** In addition to instructions, strategy definitions, and examples, we provided LLMs with two previous statements from the conversation; one from the client and one from the therapist. Using this context, we evaluated how LLMs perform when conversation history is added.

Labeling with LLMs followed the same approach as with the annotators. In each prompt, we asked the models to select one strategy from the PS Core and Facilitative categories or return ‘None’ if applicable. Prompts are available in Appendix C. We ran each LLM five times to test their stability for annotation. To evaluate the consistency of the annotation, we calculated the entropy of the labels generated by the LLMs (Li et al., 2023b). Entropy (Eq. (1)) measures the degree of uncertainty or disorder within a dataset. The higher the entropy value, the more uncertain the LLMs’ responses are.

$$u_i = - \sum_{j=1}^k P_{\theta}(a_{ij} | p_{ij}) \ln P_{\theta}(a_{ij} | p_{ij}) \quad (1)$$

where  $P_{\theta}(a_{ij} | p_{ij})$  is calculated as the frequency of one prediction among all predictions.

### 3.3 Transformer Fine-Tuning

In addition to using LLM-based models, we adopted a student-teacher approach to train four smaller transformer-based models, DeBERTa (He et al., 2020), MentalBERT (Ji et al., 2022), FlanT5 (Chung et al., 2024), and ModernBERT (Warner et al., 2024), for classifying therapist utterances. Following recent work on leveraging LLM-generated pseudo-labels for model supervision (Li et al., 2023b; Gilardi et al., 2023; Goel et al., 2023), we used annotations produced by our best-performing LLM as teacher-generated labels. From all the utterances labeled by LLM, we randomly sampled 5,000 (excluding the evaluation set) to fine-tune the transformer models. Fine-tuning was conducted separately for each annota-

tion dimension: PS Core strategies and Facilitative strategies. We evaluated all models on the human-annotated test set. Details of our model configurations and hyperparameter settings can be found in Appendix B. This setup enables scalable training while preserving privacy and maintaining alignment with domain-specific strategy definitions.

### 3.4 Therapeutic Dynamics

The previous section focused on identifying *which* strategies are used in PST. Here, we turn our attention to *how* these strategies are applied by analyzing autonomy, self-disclosure, and metaphor usage. These three dynamics were selected based on their foundational roles in therapeutic communication, autonomy as a core component of client-centered care (Entwistle et al., 2010), self-disclosure as a tool for building therapeutic alliance (Henretty and Levitt, 2010), and metaphor as a mechanism for enhancing communication and shared understanding (Stine, 2005). In consultation with the clinical expert in our team and drawing from established PST manuals (Nezu et al., 2012), we identified these dimensions as particularly informative for characterizing therapist-client interactions.

**Autonomy** is a cornerstone of effective counseling, enhancing therapeutic alliances and treatment outcomes (Dwyer et al., 2011). While its importance is well-established, measuring autonomy has largely focused on theory and behavior, overlooking linguistic markers. Previous research (Shah et al., 2022) has often relied on open-ended versus close-ended questions as a proxy for assessing autonomy. However, we adopt a more comprehensive definition grounded in the literature to better capture the nuances of autonomy in therapeutic interactions. To address this, we define autonomy through three categories: directive, non-directive, (Overholser, 1987) and N/A. Directive approaches involve therapist-led guidance, useful but potentially fostering dependency. Non-directives, such as Socratic questioning, encourage independence and self-reliance. The N/A category covers utterances unrelated to autonomy.

**Self-disclosure** in therapy involves a therapist sharing personal experiences or feelings to build trust (Marais and McBeath, 2021), validate emotions, or model behaviors, enabling a more equitable power relationship in therapy and normalizing clients’ experiences and distress (Jolley, 2019). It can be categorized as immediate, involving real-time feelings, or non-immediate, sharing past experiences

and personal anecdotes (Farber, 2006). Our framework also includes a None category for utterances lacking self-disclosure.

**Metaphor** is a symbolic approach that implies similarity between experiences, thoughts, emotions, actions, or objects (Glucksberg and Keysar, 1993). Given their crucial role in fostering learning and comprehension, counselors must be highly skilled at using them effectively in therapy (Wagener, 2017). While some studies show that LLMs struggle with metaphor and analogy comprehension (Czinczoll et al., 2022; Tong et al., 2024), we incorporated key concepts from the Conceptual Metaphor Theory (Tian et al., 2024), which has shown to enhance LLMs performance in metaphor understanding. We designed our prompts to identify metaphors, along with their corresponding source and target domains.

We prompted GPT-4o to analyze therapeutic dynamics in each therapist utterance (see Appendix C.2). To assess the quality of responses, two authors independently evaluated a random sample of 200 GPT-4o outputs, verifying their accuracy against the corresponding utterances through manual fact-checking.

### 3.5 Linguistic Patterns

We used Linguistic Inquiry and Word Count (LIWC) (Boyd et al., 2022) to examine how therapists’s language represents aspects such as emotional tone, cognitive processes, and time orientation. LIWC was chosen for its reliability in identifying psychologically meaningful language patterns, making it well-suited for analyzing subtle aspects of therapeutic communication (Lord et al., 2015).

## 4 Result

### 4.1 Performance of LLMs in Annotation

Table 1 shows the weighted F1 scores of four LLMs, including two proprietary models (GPT-4 and GPT-4o) and two open-source models (Llama-3.1 and Yi-1.5). We evaluate the models’ performance using our evaluation set (500 annotated utterances) for both ‘No context’ and ‘With context’ conditions. GPT-4o achieved the highest F1 score (0.76) with ‘No context’, outperforming all other models. GPT-4 followed closely with an F1 score of 0.68 (See Appendix D.1 for more details). All models showed lower F1 scores when provided with additional context (i.e., the two previous utterances in the conversation), and our error analysis

Model	No context			With context		
	FAC	PS	Overall	FAC	PS	Overall
GPT-4	0.65	0.76	0.68	<b>0.65</b>	0.70	0.64
GPT-4o	<b>0.78</b>	<b>0.81</b>	<b>0.76</b>	0.59	<b>0.76</b>	<b>0.66</b>
Llama-3.1	0.66	0.70	0.61	0.55	0.65	0.50
Yi-1.5	0.64	0.59	0.56	0.56	0.54	0.49

Table 1: Average weighted F1 scores for PS Core (PS), Facilitators (FAC), and overall across different LLMs. The overall score is the weighted average of all strategies’ F1 score plus None category.

Model	Mean Entropy	Std Entropy
GPT-4	0.042	0.150
GPT-4o	0.035	0.130
Llama	0.016	0.091
Yi-1.5	0.000	0.000

Table 2: Comparison of Mean and Standard Deviation of Entropy across different LLMs.

showed that this was because, despite explicitly instructing the LLMs in our prompt to focus only on the therapist’s most recent utterance and use the two previous utterances solely as context, the LLMs occasionally returned labels for the context rather than the target utterance.

Table 2 shows entropy results across five runs per LLM. Since values were identical for ‘No context’ and ‘With context’, we report a single set, with all models showing sufficiently low entropy for consistency. GPT-4 and GPT-4o had higher mean entropy than Llama-3.1 and Yi-1.5, with GPT-4o slightly more stable (0.035 vs. 0.042). Despite higher variability, GPT models’ superior F1 scores suggest this had no impact on performance. In contrast, Llama-3.1 and Yi-1.5 exhibited lower entropy, with Yi-1.5 at 0.00, potentially indicating overconfidence and poor generalization, reflected in its lower F1 score.

Overall, the results show that GPT-4o performed best without any context. Precision, recall, and F1 scores for each PST strategy are shown in Table 3. The model performed well in classifying PS Core strategies, with balanced scores across strategies like ‘Problem-Solving Positive Mindset’ and ‘Defining Problems and Goals’. For Facilitators, it was successful in identifying ‘Therapeutic Engagement’ and ‘Problem-Solving Test Review,’ while ‘Social Courtesies’ and ‘Session Management’ showed slightly lower precision. We selected GPT-4o to annotate the entire dataset of therapist utterances. Given its low mean entropy (0.035), we only ran the model once for full annotation.

	Strategies	Precision	Recall	F1-score	#Support
PS Core	Problem-solving Positive Mindset	0.85	0.85	0.85	39
	Defining Problems and Goals	0.86	0.73	0.79	66
	Generating Alternative Solutions	0.83	0.76	0.79	36
	Outcome Prediction and Planning	0.86	0.82	0.84	33
	Trying Out Solution Plan	0.67	1	0.80	18
Facilitators	Social Courtesies	0.59	0.94	0.73	25
	Session Management	0.58	0.81	0.68	64
	Therapeutic Engagement	0.69	0.95	0.80	130
	Test Review	1	0.86	0.92	24
	None	0.88	0.48	0.62	90
	<b>weighted average</b>	<b>0.77</b>	<b>0.79</b>	<b>0.76</b>	<b>525</b>

Table 3: Precision, Recall, and F1-Score for Facilitative and PS Strategies in GPT-4o (without context). Total #Support exceeds 500 (utterances in our evaluation set) due to the multi-label setup, where some utterances were labeled for both a PS Core and a Facilitative strategy.

Model	PS	FAC	Overall
DeBERTa	0.77	<b>0.73</b>	0.68
MentalBERT	0.78	0.69	0.68
ModernBERT	<b>0.8</b>	0.7	<b>0.71</b>
FLAN-T5	0.68	0.69	0.67

Table 4: Average weighted F1 scores for PS Core (PS), Facilitators (FAC), and Overall across models.

## 4.2 Transformer Fine-Tuning

After annotating the entire dataset of 14,417 therapist utterances with GPT-4o, we randomly sampled 5,000 utterances and fine-tuned four transformed-based models. Table 4 shows the average weighted F1 scores for DeBERTa, MentalBERT, FLAN-T5, and ModernBERT on PS Core and Facilitative strategies. ModernBERT achieved the highest F1 score (0.8) for PS Core strategies, while DeBERTa performed best on Facilitators (0.73). These results highlight DeBERTa’s strength in Facilitators annotation and ModernBERT’s advantage in PS Core (See Appendix D.2).

## 4.3 Therapeutic Dynamics

For autonomy, the relationship between directive and non-directive language and question types shows that 94.49% of directive utterances are closed-ended. While 59.25% of non-directive utterances are open-ended, a notable 40.75% are unexpectedly closed-ended (See the co-occurrence matrix in Appendix E). Figure 2 shows the distribution of directive and non-directive autonomy in PS Core strategies. Non-directive language dominates the early stages (‘PS Positive Mindset’ and ‘Defining Problems and Goals’), while the later stages (‘Outcome Prediction and Planning’ and ‘Trying

Out Solution Plan’) shift toward directive language. Similarly, for Facilitators (Figure 3) more structured strategies are directive while the other two show more non-directive language.

Reviewing the results of self-disclosure showed that, only 257 instances were labeled as non-immediate self-disclosure, suggesting that such disclosures are carefully considered rather than spontaneous. Furthermore, Figures 4 and 5 show metaphor usage across PS Core and Facilitative Strategies. Metaphors are most prevalent in Therapeutic Engagement (72.18%). In PS Core Strategies, they are commonly used in ‘Defining Problems and Goals’ (44.90%) and ‘Generating Alternative Solutions’ (23.13%).

## 4.4 Linguistic Patterns

Table 5 shows the distribution of all therapeutic strategies in our data with the top bigrams and LIWC categories representing each. Among PS Core strategies, ‘Defining Problems and Goals’ is most prevalent (33.48%), emphasizing therapists’ role in helping clients clarify objectives, with bigrams like “defining problem” and “solving important”. LIWC’s *reward* feature suggests clients are being motivated by potential benefits. ‘Generating Alternative Solutions’ is another significant strategy (20.03%), with words such as “brainstorming ideas” and “strategies tactics” showing the focus on exploring different options to address issues. The LIWC indicators *insight* and *curiosity* align with this focus on creativity and open-mindedness. Similarly, ‘Problem-Solving Positive Mindset’ includes “positive attitude” and “healthy thinking”, highlighting therapists’ encouragement toward optimism and resilience, supported by LIWC features

	Strategies	Top Bigrams	Top LIWC	Distribution
PS Core	Defining Problems and Goals	(goals, around), (defining, problem), (facts, around), (obstacles, goals), (clear, language), (meeting, goals), (solving, important), (one, goals)	reward	33.48%(1,743)
	Generating Alternative Solutions	(brainstorming, ideas), (strategies, tactics), (defer, judgment), (leads, quality), (quantity, leads), (three, things), (variety, ideas), (lots, ideas), (good, thing), (big, idea)	insight, curiosity	20.03%(1,043)
	Problem-solving	(strong, emotion), (way, feel), (affects, way), (healthy, thinking), (make, mistakes), (positive, attitude), (things, change), (thinking, rules), (really, helpful), (use, feelings), (try, balance), (problems, challenges), (feelings, adaptively)	emotion, feeling	11.75%(612)
	Outcome Prediction and Planning	(social, consequences), (personal, consequences), (positive, negative), (solve, problem), (pros, cons), (last, time), (personal, social), (rough, screening), (would, affect)	risk, focusfuture	10.18%(530)
	Trying Out Solution Plan	(positive, negative), (see, works), (make, sense), (really, hard), (go, back), (solve, problem), (really, important), (effort, put), (get, support), (extra, support), (talk, trying)	focuspast, focusfuture	2.97%(155)
Facilitators	Therapeutic Engagement	(really, good), (makes, sense), (right, right), (last, time), (oh, gosh), (good, good), (make, sure), (might, look), (one, things), (goals, around)	focuspresent	52.19%(5,918)
	Session Management	(next, time), (time, meet), (last, time), (defining, problem), (alternative, solutions), (next, week), (let, see), (going, talk), (time, talk), (brainstorming, ideas)	focusfuture	15.80%(1,792)
	Social Courtesies	(thank, much), (look, forward), (thank, time), (talk, next), (taking, time), (next, time), (good, good), (thank, appreciate), (take, care), (really, helpful)	emotion, prosocial	8.69%(986)
	Test Review	(problem, solving), (positive, rational), (impulsive, avoidant), (positive, negative), (high, positive), (rational, impulsive), (negative, orientation), (negative, rational), (approach, problem), (one, questionnaires), (positive, orientation), (five, dimensions)	insight	1.72%(196)
	None	-	-	21.57%(3,111)

Table 5: Overview of strategies and their most frequent bigrams and LIWC features with distribution. (Note that the distribution percentages do not sum to 100% because the classification was multi-label: a single utterance could be assigned both a PS Core and a Facilitative strategy.)



Figure 2: Distribution of Autonomy in PS Core




Figure 3: Distribution of Autonomy in Facilitators

of *emotion* and *feeling*. Outcome Prediction and Planning uses “pros cons” and “solve problem”, emphasizing consequence evaluation. The prevalence of LIWC’s *risk* and *focus future* reinforces a forward-thinking approach in decision-making.

On the Facilitators side, the bigrams highlight how therapists build rapport and guide session flow. In ‘Therapeutic Engagement’, phrases like “really good” use affirming language to foster a supportive, collaborative atmosphere, helping clients feel understood. ‘Social Courtesies,’ with phrases like ‘thank much’ and ‘look forward,’ add warmth to interactions, maintaining a respectful and friendly tone. The top LIWC features, *emotion* and *prosocial*, reflect the positive social behaviors therapists use to foster mutual respect and connection. Meanwhile, in ‘Test Review’, more formal words like “positive rational” and “impulsive avoidant” suggest an analytical approach, where therapists review the client’s PST test results in a structured, evaluative manner (See Appendix F for more information).

#### 4.5 Strategy Progression in Therapy Sessions.

We analyzed how strategies progress across three different visits of caregivers in our dataset. In Figure 6, early sessions focus on enhancing ‘Problem-

Solving Mindset’ (Step One) and ‘Defining Problems and Goals’ (Step Two), indicating that the initial phase of therapy is centered on improving caregivers’ mindsets and helping them define the challenges they face. As therapy progresses into the second and third visits, there is a shift towards more advanced strategies such as ‘Generating Alternative Solutions’ (Step Three) and ‘Outcome Prediction and Planning’ (Step Four), showing that the focus moves towards exploring possible solutions and planning for their implementation. In the final phase (Visit Three), the ‘Trying Out Solution Plan’ (Step Five) becomes prominent, indicating a transition to applying and testing solutions in real life. Figure 7 shows the consistent use of Facilitative strategies like ‘Therapeutic Engagement,’ ‘Session Management,’ and ‘Social Courtesies,’ reinforcing their essential role throughout therapy. However, ‘Test Review’ is more frequent in Visit One, reflecting the therapist’s focus on reviewing the client’s test results taken prior to starting the session before progressing to intervention. This early emphasis lays the groundwork for later sessions. These patterns further validate the effectiveness of our classifier.



Figure 4: Proportion of Metaphors in PS Core




Figure 5: Proportion of Metaphors in Facilitators

## 5 Discussion

**Model Performance and the Need for Open-Source Improvements.** Our results show that GPT-4o outperformed GPT-4 and open-source models in annotating therapeutic strategies, achieving an F1 score of 0.76 without supplemental context. While Yi had the lowest entropy, GPT-4o showed greater reliability, highlighting the trade-off between performance stability and model reliability in mental health applications. Despite GPT-4o’s success, improving open-source models is crucial for enhancing accuracy, accessibility, and competition in healthcare AI. From an ethical perspective, open-source models promote transparency and accountability, ensuring that healthcare AI advancements benefit a broader community and expand access to high-quality mental health resources, unlike proprietary models, which can limit access and obscure critical decision-making processes (Welsh et al., 2024).

**Improving LLMs for Emotional and Social Understanding.** GPT-4o excelled in labeling structured, goal-oriented PS core strategies, achieving high precision and recall. This highlights LLMs’ strength in processing clear, task-oriented language. However, GPT-4o struggled with Facilitative strategies such as ‘Session Management’ and ‘Social Courtesies’, which rely on emotional and social cues. This underscores a key limitation: while LLMs handle problem-solving tasks well, they struggle with the nuances of interpersonal communication, lacking the ability to interpret human emotions accurately. These challenges arise from implicit, context-dependent interactions, which LLMs often misinterpret, sometimes displaying an over-realistic sense of empathy (Sorin et al., 2023; Roshanaei et al., 2024). To improve effectiveness in therapeutic settings, future work should refine LLMs using expert-in-the-loop systems, allowing for real-time human feedback to enhance emotional and social cue recognition (Afzal et al., 2024; Sun et al., 2023).

**Efficient and Privacy-Conscious Therapy Anno-**

**tation.** Our fine-tuned transformer-based models show the effectiveness of the distillation (student-teacher) method, achieving strong performance. ModernBERT outperformed our second-best LLM, GPT-4, with an overall F1 score of 0.71. It achieved 0.8 F1 in PS Core strategies, while DeBERTa specialized in Facilitative strategies, reaching an F1 score of 0.73. This suggests that smaller models can effectively annotate distinct therapeutic strategies, offering a scalable and efficient alternative for therapy content analysis. Additionally, these models provide a privacy-conscious solution for processing sensitive therapy data, ensuring confidentiality and trust in digital therapeutic environments (Shah et al., 2022).

**How PST is Delivered.** Beyond classifying strategies, understanding how therapists apply therapeutic tools like autonomy, self-disclosure, and metaphor offers deeper insight into PST’s communication dynamics. In early PS Core stages, therapists favor non-directive language to foster client exploration, shifting to directive language later for clarity and structure, aligning with a client-centered problem-solving approach (Nezu et al., 2012). Examining question types and autonomy, we found that some closed-ended questions still encouraged elaboration, which our domain expert termed “imposters.” For example, “Do you have any questions about what we’ve gone over or for this process?”, though closed-ended, it invites further discussion, blurring the line between directive and non-directive communication. This suggests that therapists sometimes frame questions strategically to encourage deeper engagement, even when using traditionally restrictive formats.

Therapists also use non-immediate self-disclosure strategically and selectively, as reflected in the low frequency of such utterances in our dataset. In Facilitators, self-disclosures often involve personal anecdotes to build rapport, whereas, in PS Core, they serve a more functional role, modeling problem-solving strategies for clients. This finding is supported by previous work (Knox






Figure 6: Distribution of PS Core in three visits

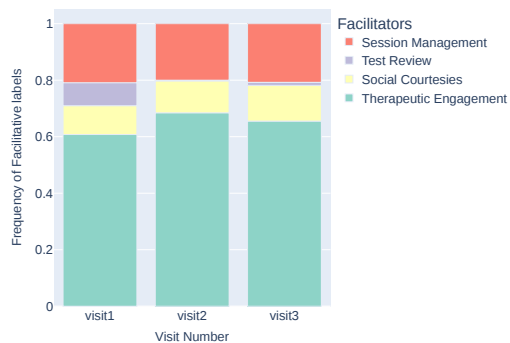


Figure 7: Distribution of Facilitators in three visits

and Hill, 2003) which identified similar patterns of strategic disclosure timing and emphasized how self-disclosure can serve as different functions at different stages of therapy. Metaphors help illustrate ideas and enhance emotional connection in Therapeutic Engagement (Lyddon et al., 2001; Tay, 2014), while in PS Core strategies, they help conceptualization and creativity. However, they decline in action-oriented phases, where direct communication becomes essential. Therapists adapt their language to session goals, fostering exploration early on and shifting to structured, goal-driven communication as therapy progresses.

**Broader Implications.** Our model has meaningful real-world applications for PST practice and documentation. It can assist in annotating and analyzing therapy sessions, helping document which PST strategies were employed and how effectively they were implemented, and serving valuable educational purposes for therapists-in-training. As LLM-assisted therapy models become more prevalent, our work provides important guidance for human-AI collaborative approaches. By identifying which PST strategies are most crucial and how they should be properly sequenced for building client rapport, we can better inform the development of LLM-assisted therapy systems where human therapists remain in the loop. This allows for more effective integration of AI support tools while maintaining the essential human element of therapeutic relationships. The model can also serve as an evaluation framework for assessing whether AI-assisted therapy is appropriately implementing PST strategies and maintaining therapeutic alliances through their interactions.

## 6 Conclusion

This study highlights the potential of LLMs like GPT-4o for effective annotation of therapeutic di-

alogues. GPT-4o achieved a notable F1 score of 0.76, with a low entropy value of 0.035, demonstrating its efficiency and reliability in categorizing PST strategies. We fine-tuned transformer-based models on a sample of 5,000 therapist utterances annotated by GPT-4o and showed that ModernBERT exceeded the performance of our second-best LLM, GPT-4, showcasing the effectiveness of using LLM-generated annotations to fine-tune smaller, domain-specific models. We also expanded the PST codebook to include a new Facilitators dimension, enriching our analysis of therapist-client interactions. We also analyzed Therapeutic Dynamics (autonomy, self-disclosure, and metaphor) and linguistic patterns in our data to better understand ‘how’ PST strategies are used and in what contexts. Our findings showed that all models performed better at identifying ‘PS Core’ strategies, indicating a strong capability in recognizing structured therapeutic elements. Our research offers valuable insights into enhancing technology integration in mental health care, aiming to make therapeutic interventions both more accessible and more effective.

## 7 Limitations

Our study has several limitations, including the dependence on ‘transcribed’ therapy sessions, which may not fully capture the dynamic aspects of live interactions such as tone and non-verbal cues. The use of LLMs like GPT-4o and GPT-4 introduces potential biases from their internet-based training data, impacting the interpretation of therapeutic strategies. Furthermore, our analysis is solely based on English-language data, which limits the model’s applicability in non-English-speaking contexts. These limitations highlight the technology’s current stage and its focus on textual analysis, excluding the broader emotional and psychological

facets of therapy sessions.

In addition, we only focus on PST-related framework and analysis which may not be generalizable to other therapeutic methods like cognitive-behavioral therapy (CBT) or interpersonal therapy (IPT). These methods may involve different communication styles and dynamics that the LLMs we used may not accurately interpret. Future work will analyze the adaptability of our model for broader therapy analysis. Additionally, cultural and contextual differences in therapy could further limit the model's applicability across diverse mental health settings. Future research could address this by testing the models in a broader range of therapeutic approaches and contexts.

Finally, we did not compare our results with any LLM-generated therapy data, which could offer a valuable comparison between real-world therapeutic interactions and simulated data produced by LLMs. Our focus is exclusively on the characteristics of authentic, real-world data from therapy sessions. Future research could explore this comparison, examining how well LLM-generated interactions align with or diverge from real-world therapeutic PST dialogues, potentially offering insights into improving the practicality and effectiveness of AI-driven interventions in mental health care.

## 8 Ethical Considerations

We are collaborating with a team of PST experts who were in charge of collecting and transcribing the data. We only have access to secondary data, and the primary data was collected as part of a larger task approved by the primary institution's review board (Demiris et al., 2019). As a result, we do not have access to any identifiable information about the clients or therapists. We understand that using anonymized therapy data raises concerns regarding privacy and the potential for re-identification of people involved in these sessions. We ensured that all data used in our analysis was thoroughly cleaned, with the help of our collaborators, and that strict measures were put in place to remove any remaining information from the transcriptions.

The application of LLMs in healthcare, particularly in sensitive areas like therapy, necessitates careful consideration of their implications. Ensuring responsible development and use of AI technologies in mental health care is crucial, requiring vigilance for biases, stakeholder engagement, and

strong safeguards to maintain patient confidentiality and the integrity of the therapeutic process. This paper does not endorse the use of LLMs in therapeutic settings, nor does it claim that they are ready for such applications. Rather, our goal is to systematically evaluate the behavior of current LLMs when applied in therapeutic contexts. The use of AI in mental health care raises ethical concerns about automating therapy. AI tools, like LLMs, should assist rather than replace therapists, ensuring they do not undermine the personal, empathetic nature of therapy, which is crucial for effective treatment. While AI can provide useful insights, therapists must remain the primary decision-makers, as AI cannot (yet) fully grasp the complexities of human behavior and emotions as shown in our results and other related work. Accountability should always rest with human professionals to preserve the integrity and effectiveness of mental health care.

## 9 Acknowledgment

We are deeply grateful to Dr. Bum Chul Kwon, Max Song, and DongWhan Jun for their invaluable insights and assistance throughout this project. This paper is based upon work supported by the National Science Foundation under Grant #2144880. The intervention that generated the therapy transcripts used in this study was supported in part by the National Institute of Nursing Research (NINR) of the National Institutes of Health under award number 2R01NR012213.

## References

- Anum Afzal, Alexander Kowsik, Rajna Fani, and Florian Matthes. 2024. [Towards optimizing and evaluating a retrieval augmented QA chatbot using LLMs with human-in-the-loop](#). In *Proceedings of the Fifth Workshop on Data Science with Human-in-the-Loop (DaSH 2024)*, pages 4–16, Mexico City, Mexico. Association for Computational Linguistics.
- Laberiano Andrade-Arenas, Cesar Yactayo-Arias, and Félix Pucuhuayla-Revatta. 2024. [Therapy and emotional support through a chatbot](#). *International Journal of Online and Biomedical Engineering (iJOE)*, 20(02):pp. 114–130.
- American Psychological Association APA. 2023. [Stress in america 2023: A nation recovering from collective trauma](#).
- Patricia Arean, Mark Hegel, Steven Vannoy, Ming-Yu Fan, and Jurgen Unutzer. 2008. Effectiveness of problem-solving therapy for older, primary care

- patients with depression: results from the impact project. *The Gerontologist*, 48(3):311–323.
- Eileen Bendig, Benjamin Erb, Lea Schulze-Thuesing, and Harald Baumeister. 2022. The next generation: chatbots in clinical psychology and psychotherapy to foster mental health—a scoping review. *Verhaltens-therapie*, 32(Suppl. 1):64–76.
- Layla Bouzoubaa, Elham Aghakhani, and Rezvaneh Rezapour. 2024a. [Words matter: Reducing stigma in online conversations about substance use with large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9139–9156, Miami, Florida, USA. Association for Computational Linguistics.
- Layla Bouzoubaa, Elham Aghakhani, Max Song, Quang Trinh, and Shadi Rezapour. 2024b. [Decoding the narratives: Analyzing personal drug experiences shared on Reddit](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 6131–6148, Bangkok, Thailand. Association for Computational Linguistics.
- Ryan L Boyd, Ashwini Ashokkumar, Sarah Seraj, and James W Pennebaker. 2022. The development and psychometric properties of liwc-22. *Austin, TX: University of Texas at Austin*, 10.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Matusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.
- Yu Ying Chiu, Ashish Sharma, Inna Wanyin Lin, and Tim Althoff. 2024. A computational framework for behavioral assessment of llm therapists. *arXiv preprint arXiv:2401.00820*.
- Young Min Cho, Sunny Rai, Lyle Ungar, João Sedoc, and Sharath Guntuku. 2023. [An integrative survey on mental health conversational agents to bridge computer science and medical perspectives](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11346–11369, Singapore. Association for Computational Linguistics.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.
- Tamara Czinczoll, Helen Yannakoudakis, Pushkar Mishra, and Ekaterina Shutova. 2022. [Scientific and creative analogies in pretrained language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2094–2100, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- George Demiris, Debra Parker Oliver, Karla Washington, and Kenneth Pike. 2019. A problem-solving intervention for hospice family caregivers: a randomized clinical trial. *Journal of the American Geriatrics Society*, 67(7):1345–1352.
- Dorottya Demszky, Diyi Yang, David S. Yeager, Christopher J. Bryan, Margaret Clapper, Susannah Chandhok, Johannes C. Eichstaedt, Cameron A. Hecht, Jeremy P. Jamieson, Meghann Johnson, Michaela Jones, Danielle Krettek-Cobb, Leslie Lai, Nirel JonesMitchell, Desmond C. Ong, Carol S. Dweck, James J. Gross, and James W. Pennebaker. 2023. [Using large language models in psychology](#). *Nature Reviews Psychology*, 2:688 – 701.
- Lynndall A Dwyer, Matthew J Hornsey, Laura GE Smith, Tian PS Oei, and Genevieve A Dingle. 2011. Participant autonomy in cognitive behavioral group therapy: An integration of self-determination and cognitive behavioral theories. *Journal of Social and Clinical Psychology*, 30(1):24–46.
- Vikki A Entwistle, Stacy M Carter, Alan Cribb, and Kirsten McCaffery. 2010. Supporting patient autonomy: the importance of clinician-patient relationships. *Journal of general internal medicine*, 25:741–745.
- Barry Alan Farber. 2006. *Self-disclosure in psychotherapy*. Guilford Press.
- Daniil Filienko, Yinzhou Wang, Caroline El Jazmi, Serena Xie, Trevor Cohen, Martine De Cock, and Weichao Yuwen. 2024. Toward large language models as a therapeutic tool: Comparing prompting techniques to improve gpt-delivered problem-solving therapy. *arXiv preprint arXiv:2409.00112*.
- Zvi D Gellis and Bonnie Kenaley. 2008. Problem-solving therapy for depression in adults: a systematic review. *Research on social work practice*, 18(2):117–131.
- Hossein Ghamari Kivi and Fatemeh Jamshid-doust Miyanroudi. 2022. Tracking sequences of patient-therapist dialogues by coding responses during integrative psychotherapy.
- Fabrizio Gildardi, Meysam Alizadeh, and Maël Kubli. 2023. Chatgpt outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30):e2305016120.
- Sam Glucksberg and Boaz Keysar. 1993. How metaphors work. *Metaphor and thought*, 2:401–424.
- Akshay Goel, Almog Gueta, Omry Gilon, Chang Liu, Sofia Erell, Lan Huong Nguyen, Xiaohong Hao, Bolous Jaber, Shashir Reddy, Rupesh Kartha, et al.

2023. Llms accelerate annotation for medical information extraction. In *Machine Learning for Health (ML4H)*, pages 82–100. PMLR.
- Keith Harrigan, Carlos Aguirre, and Mark Dredze. 2020. Do models of mental health based on social media data generalize? In *Findings of the association for computational linguistics: EMNLP 2020*, pages 3774–3788.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Xingwei He, Zhenghao Lin, Yeyun Gong, A-Long Jin, Hang Zhang, Chen Lin, Jian Jiao, Siu Ming Yiu, Nan Duan, and Weizhu Chen. 2024. [AnnoLLM: Making large language models to be better crowdsourced annotators](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 6: Industry Track)*, pages 165–190, Mexico City, Mexico. Association for Computational Linguistics.
- Jennifer R Henretty and Heidi M Levitt. 2010. The role of therapist self-disclosure in psychotherapy: A qualitative review. *Clinical psychology review*, 30(1):63–77.
- Emma Hoes, Sacha Altay, and Juan Bermeo. 2023. Leveraging chatgpt for efficient fact-checking. 2023. Available on: <https://doi.org/10.31234/osf.io/qnjkf> and.
- Yining Hua, Fenglin Liu, Kailai Yang, Zehan Li, Yi-han Sheu, Peilin Zhou, Lauren V Moran, Sophia Ananiadou, and Andrew Beam. 2024. Large language models in mental health care: a scoping review. *arXiv preprint arXiv:2401.02984*.
- Saadat Izadi and Mohamad Forouzanfar. 2024. [Error correction and adaptation in conversational ai: A review of techniques and applications in chatbots](#). *AI*, 5(2):803–841.
- Shaoxiong Ji, Tianlin Zhang, Luna Ansari, Jie Fu, Prayag Tiwari, and Erik Cambria. 2022. [MentalBERT: Publicly available pretrained language models for mental healthcare](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 7184–7190, Marseille, France. European Language Resources Association.
- Helen Kathryn Jolley. 2019. I’m human too: Person-centred counsellors’ lived experiences of therapist self-disclosure. *European Journal for Qualitative Research in Psychotherapy*, 9:12–26.
- Hannah Kim, Kushan Mitra, Rafael Li Chen, Sajjadur Rahman, and Dan Zhang. 2024. [MEGAnno+: A human-LLM collaborative annotation system](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 168–176, St. Julians, Malta. Association for Computational Linguistics.
- Sarah Knox and Clara E Hill. 2003. Therapist self-disclosure: Research-based suggestions for practitioners. *Journal of clinical psychology*, 59(5):529–539.
- Harsh Kumar, Yiyi Wang, Jiakai Shi, Ilya Musabirov, Norman AS Farb, and Joseph Jay Williams. 2023. Exploring the use of large language models for improving the awareness of mindfulness. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–7.
- Taja Kuzman, Igor Mozetič, and Nikola Ljubešić. 2023. [Chatgpt: Beginning of an end of manual linguistic data annotation? use case of automatic genre identification](#). *Preprint*, arXiv:2303.03953.
- Hannah R Lawrence, Renee A Schneider, Susan B Rubin, Maja J Matarić, Daniel J McDuff, and Megan Jones Bell. 2024. The opportunities and risks of large language models in mental health. *JMIR Mental Health*, 11(1):e59479.
- Younghun Lee, Dan Goldwasser, and Laura Schwab Reese. 2024. [Towards understanding counseling conversations: Domain knowledge and large language models](#). In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 2032–2047, St. Julian’s, Malta. Association for Computational Linguistics.
- Anqi Li, Lizhi Ma, Yaling Mei, Hongliang He, Shuai Zhang, Huachuan Qiu, and Zhenzhong Lan. 2023a. Understanding client reactions in online mental health counseling. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10358–10376.
- Jiyi Li. 2024. A comparative study on annotation quality of crowdsourcing and llm via label aggregation. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6525–6529. IEEE.
- Minzhi Li, Taiwei Shi, Caleb Ziemis, Min-Yen Kan, Nancy Chen, Zhengyuan Liu, and Diyi Yang. 2023b. [CoAnnotating: Uncertainty-guided work allocation between human and large language models for data annotation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1487–1505, Singapore. Association for Computational Linguistics.
- Baihan Lin, Djallel Bouneffouf, Yulia Landa, Rachel Jespersen, Cheryl Corcoran, and Guillermo Cecchi. 2025. Compass: Computational mapping of patient-therapist alliance strategies with language modeling. *Translational Psychiatry*, 15(1):1–15.
- Yilun Liu, Shimin Tao, Weibin Meng, Feiyu Yao, Xiaofeng Zhao, and Hao Yang. 2024. Logprompt:

- Prompt engineering towards zero-shot and interpretable log analysis. In *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings*, pages 364–365.
- Sarah Peregrine Lord, Elisa Sheng, Zac E Imel, John Baer, and David C Atkins. 2015. More than reflections: Empathy in motivational interviewing includes language style synchrony between therapist and client. *Behavior therapy*, 46(3):296–303.
- William J Lyddon, Alison L Clay, and Cheri L Sparks. 2001. Metaphor and change in counseling. *Journal of Counseling & Development*, 79(3):269–274.
- Zilin Ma, Yiyang Mei, and Zhaoyuan Su. 2023. Understanding the benefits and challenges of using large language model-based conversational agents for mental well-being support. In *AMIA Annual Symposium Proceedings*, volume 2023, page 1105. American Medical Informatics Association.
- John M Malouff, Einar B Thorsteinsson, and Nicola S Schutte. 2007. The efficacy of problem solving therapy in reducing mental and physical health problems: a meta-analysis. *Clinical psychology review*, 27(1):46–57.
- Genevieve Marais and Alistair McBeath. 2021. Therapists’ lived experience of self-disclosure. *European Journal for Qualitative Research in Psychotherapy*, 11:72–86.
- Matthew Louis Mauriello, Thierry Lincoln, Grace Hon, Dorien Simon, Dan Jurafsky, and Pablo Paredes. 2021. Sad: A stress annotated dataset for recognizing everyday stressors in sms-like conversational systems. In *Extended abstracts of the 2021 CHI conference on human factors in computing systems*, pages 1–7.
- Maitrey Mehta, Derek Caperton, Katherine Axford, Lauren Weitzman, David Atkins, Vivek Srikumar, and Zac Imel. 2022. Psychotherapy is not one thing: simultaneous modeling of different therapeutic approaches. In *Proceedings of the Eighth Workshop on Computational Linguistics and Clinical Psychology*, pages 47–58.
- Meta. Meta llama 3.1. <https://ai.meta.com/blog/meta-llama-3-1/>. Accessed: 2024-10-10.
- Arthur M Nezu, Christine M Nezu, and Thomas J D’Zurilla. 2012. *Problem-solving therapy: A treatment manual*. springer publishing company.
- OpenAI. Gpt-4 model documentation. <https://platform.openai.com/docs/models/gpt-4o>. Accessed: 2024-10-10.
- OpenAI. 2024. GPT-4, Version 0125 Preview. Accessed: 2024-10-10.
- James C Overholser. 1987. Facilitating autonomy in passive-dependent persons: An integrative model. *Journal of Contemporary Psychotherapy*, 17(4):250–269.
- Aria Pessianzadeh and Rezvaneh Rezapour. 2025. Exploring stance on affirmative action through reddit narratives. In *Proceedings of the 17th ACM Web Science Conference 2025*, pages 52–63.
- Kelly Posner, Gregory K Brown, Barbara Stanley, David A Brent, Kseniya V Yershova, Maria A Oquendo, Glenn W Currier, Glenn A Melvin, Laurence Greenhill, Sa Shen, et al. 2011. The columbia–suicide severity rating scale: initial validity and internal consistency findings from three multisite studies with adolescents and adults. *American journal of psychiatry*, 168(12):1266–1277.
- Laria Reynolds and Kyle McDonell. 2021. Prompt programming for large language models: Beyond the few-shot paradigm. In *Extended abstracts of the 2021 CHI conference on human factors in computing systems*, pages 1–7.
- Eugénia Ribeiro, Antonio P Ribeiro, Miguel M Gonçalves, Adam O Horvath, and William B Stiles. 2013. How collaboration in therapy becomes therapeutic: The therapeutic collaboration coding system. *Psychology and Psychotherapy: Theory, Research and Practice*, 86(3):294–314.
- Mahnaz Roshanaei, Rezvaneh Rezapour, and Magy Seif El-Nasr. 2024. Talk, listen, connect: Navigating empathy in human-ai interactions. *arXiv preprint arXiv:2409.15550*.
- Raj Sanjay Shah, Faye Holt, Shirley Anugrah Hayati, Aastha Agarwal, Yi-Chia Wang, Robert E. Kraut, and Diyi Yang. 2022. [Modeling motivational interviewing strategies on an online peer-to-peer counseling platform](#). *Proc. ACM Hum.-Comput. Interact.*, 6(CSCW2).
- Ashish Sharma, Kevin Rushton, Inna Lin, David Wadden, Khendra Lucas, Adam Miner, Theresa Nguyen, and Tim Althoff. 2023. [Cognitive reframing of negative thoughts through human-language model interaction](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9977–10000, Toronto, Canada. Association for Computational Linguistics.
- Alexander Smith, Stefanie Hachen, Roman Schleifer, Dinesh Bhugra, Anna Buadze, and Michael Liebrecht. 2023. Old dog, new tricks? exploring the potential functionalities of chatgpt in supporting educational methods in social psychiatry. *International Journal of Social Psychiatry*, 69(8):1882–1889.
- Ruijun Song and Nana Sun. 2019. The efficacy of psychological capital intervention (pci) for depression from the perspective of positive psychology: a pilot study. *Frontiers in psychology*, 10:459118.

- Vera Sorin, Danna Brin, Yiftach Barash, Eli Konen, Alexander Charney, Girish Nadkarni, and Eyal Klang. 2023. Large language models (llms) and empathy-a systematic review. *medRxiv*, pages 2023–08.
- John J Stine. 2005. The use of metaphors in the service of the therapeutic alliance and therapeutic communication. *Journal of the American Academy of Psychoanalysis and Dynamic Psychiatry*, 33(3):531–545.
- Xin Sun, Jos A. Bosch, Jan De Wit, and Emiel Kraemer. 2023. [Human-in-the-loop interaction for continuously improving generative model in conversational agent for behavioral intervention](#). In *Companion Proceedings of the 28th International Conference on Intelligent User Interfaces, IUI '23 Companion*, page 99–101, New York, NY, USA. Association for Computing Machinery.
- Yi Tang, Chia-Ming Chang, and Xi Yang. 2024. Pdfchatannotator: A human-llm collaborative multimodal data annotation tool for pdf-format catalogs. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*, pages 419–430.
- Dennis Tay. 2014. Metaphor theory for counselling professionals. *Bloomsbury companion to cognitive linguistics*, pages 352–367.
- Yuan Tian, Nan Xu, and Wenji Mao. 2024. [A theory guided scaffolding instruction framework for LLM-enabled metaphor reasoning](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7738–7755, Mexico City, Mexico. Association for Computational Linguistics.
- Xiaoyu Tong, Rochelle Choenni, Martha Lewis, and Ekaterina Shutova. 2024. [Metaphor understanding challenge dataset for LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3517–3536, Bangkok, Thailand. Association for Computational Linguistics.
- Petter Törnberg. 2023. Chatgpt-4 outperforms experts and crowd workers in annotating political twitter messages with zero-shot learning. *arXiv preprint arXiv:2304.06588*.
- Alwin E Wagener. 2017. Metaphor in professional counseling. *Professional Counselor*, 7(2):144–154.
- Hongru Wang, Huimin Wang, Lingzhi Wang, Minda Hu, Rui Wang, Boyang Xue, Hongyuan Lu, Fei Mi, and Kam-Fai Wong. 2023. Tpe: Towards better compositional reasoning over conceptual tools with multi-persona collaboration. *arXiv preprint arXiv:2309.16090*.
- Lu Wang, Munif Ishad Mujib, Jake Williams, George Demiris, and Jina Huh-Yoo. 2021a. An evaluation of generative pre-training model-based therapy chatbot for caregivers. *arXiv preprint arXiv:2107.13115*.
- Shuohang Wang, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2021b. Want to reduce labeling cost? gpt-3 can help. *arXiv preprint arXiv:2108.13487*.
- Xinru Wang, Hannah Kim, Sajjadur Rahman, Kushan Mitra, and Zhengjie Miao. 2024. Human-llm collaborative annotation through effective verification of llm labels. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–21.
- Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Nathan Cooper, Griffin Adams, Jeremy Howard, and Iacopo Poli. 2024. [Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference](#). *Preprint*, arXiv:2412.13663.
- Ceilidh Welsh, Susana Román García, Gillian C Barnett, and Raj Jena. 2024. Democratising artificial intelligence in healthcare: community-driven approaches for ethical solutions. *Future Healthcare Journal*, 11(3):100165.
- WHO. 2022. [WHO highlights urgent need to transform mental health and mental health care](#).
- Kailai Yang, Tianlin Zhang, Ziyang Kuang, Qianqian Xie, Jimin Huang, and Sophia Ananiadou. 2024. Mental-lama: interpretable mental health analysis on social media with large language models. In *Proceedings of the ACM on Web Conference 2024*, pages 4489–4500.
- Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, et al. 2024. Yi: Open foundation models by 01. ai. *arXiv preprint arXiv:2403.04652*.
- Zhonghua Zheng, Lizi Liao, Yang Deng, and Liqiang Nie. 2023. Building emotional support chatbots in the era of llms. *arXiv preprint arXiv:2308.11584*.
- Yiming Zhu, Peixian Zhang, Ehsan-Ul Haq, Pan Hui, and Gareth Tyson. 2023. Can chatgpt reproduce human-generated labels? a study of social computing tasks. *arXiv preprint arXiv:2304.10145*.
- Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. Can large language models transform computational social science? *Computational Linguistics*, 50(1):237–291.

## A Therapists Strategies

Table 6 presents the final list of therapist strategies for PST, with definitions and examples for each category.

Category	Strategy	Definition	Example
PS Core	Problem-solving Positive Mindset (Step One)	Strategies building a positive approach towards challenges, emphasizing optimism and emotional awareness.	<i>"uh, they suggest trying to think of problems as challenges rather than threats."</i>
	Defining Problems and Goals (Step Two)	Setting achievable goals and identifying factual problems while recognizing obstacles.	<i>"Are there any obstacles to those goals? Can we list anything that is getting in the way?"</i>
	Generating Alternative Solutions (Step Three)	Brainstorming solutions, differentiating between strategies, and encouraging open idea exploration.	<i>"So let's brainstorm some specific strategies and tactics for how you can make that happen."</i>
	Outcome Prediction and Planning (Step Four)	Narrowing down solutions, assessing impacts, and creating an action plan.	<i>"Working through this worksheet...short-term personal consequences of self-care are pretty positive, right?"</i>
	Trying Out Solution Plan (Step Five)	Implementing and assessing the solution strategy, troubleshooting, and acknowledging successes.	<i>"Then just monitoring your outcomes. Has the problem improved in some ways?"</i>
Facilitators	Social Courtesies*	Polite, surface-level interactions like greetings, small talk, and farewells to create a friendly atmosphere.	<i>"thank you again for your time today, and have a really good week."</i>
	Session Management*	The therapist provides guidance on session logistics, manages time, keeps discussions on track, and helps plan future sessions.	<i>"So you got Skype up and running, and that is so fun!"</i>
	Therapeutic Engagement*	The therapist shows attention, validates emotions, offers support, reflects feelings, praises strengths, and provides comfort or personal insight.	<i>"I'm glad it's starting to show some small improvements."</i>
	Test Review*	The therapist explains the PS Test and reviews the client's results across five dimensions: positive/negative problem orientation, rational, impulsive, and avoidant problem-solving styles.	<i>"So positive and rational were your highest scores. Your positive was, uh, really, is almost as high as it could possibly get."</i>

Table 6: Overview of Therapist Strategies for PST, \*denotes newly created codes

## B Model Setting

**LLM parameters:** temperature = 0;  
max-tokens = 500;  
openai-api-version='gpt-4o-2024-05-13',  
'gpt-4-0125-preview',  
'meta-llama/Llama-3.1-70B-Instruct',  
'01-ai/Yi-1.5-34B-Chat'

**Transformer-based Models Parameters:** We fine-tuned each model using the Adam optimizer with the following hyperparameters: learning\_rate=2e-5, batch\_size=8, num\_train\_epochs=3, weight\_decay=0.01. The models used were: DeBERTa:microsoft/deberta-v3-base, MentalBERT:mental/mental-bert-base-uncased, FlanT5:google/flan-t5-base, ModernBERT:answerdotai/ModernBERT-base.

## C Prompts

In this section, we present the specific prompt used to obtain annotations from LLMs (the few-shot examples are not included here).

## C.1 Classification of PST Strategies

### C.1.1 Without Context

**Instruction.** You are a highly helpful assistant tasked with annotating the therapist's behavior. Your role is to categorize the therapist's behavior using a list of strategies provided below.

Your Task: Select one strategy from the 'PS core' category and one strategy from the Facilitators category that best aligns with the therapist's utterance. If none of the strategies are applicable, return 'None' for that category.

**Definition.** Definitions of Therapist's Strategies for Problem-Solving Therapy:

PS core Strategies: 1- Problem-solving Positive Mindset: Techniques that foster a positive outlook, such as promoting optimism, reframing negative thoughts, visualizing success, practicing gratitude, mindfulness, positive self-talk, and emotional awareness. 2- Defining Problems and Goals: Identifying clear, achievable goals, recognizing obstacles, and determining which specific issues to address in therapy. 3- Generating Alternative Solutions: Brainstorming a

variety of solutions without judgment, distinguishing between strategies and tactics, and encouraging open exploration of ideas. 4- Outcome Prediction and Planning: Narrowing down solutions, evaluating their impact (personal, social, temporal), and creating a detailed action plan for implementation. 5- Trying Out Solution Plan: Implementing the solution, assessing its effectiveness, troubleshooting issues, seeking support, and acknowledging successes in the process.

Facilitative Strategies: 6- Social Courtesies: Engaging in polite, yet somewhat superficial social interactions like greetings, small talk (e.g., asking about the weather or weekend plans), thanking the client for their time, and saying goodbye. 7- Session Management: Managing session logistics, such as informing the client about session length, guiding the direction of the conversation, refocusing when off-topic, or scheduling future sessions. 8- Therapeutic Engagement: Actively listening, validating the client's feelings, providing emotional support, recalling previous information shared by the client, reflecting emotions, normalizing experiences, commenting on client strengths, offering personal disclosure, and offering comfort. 9- Test Review: Discussing the client's Problem-Solving Test results, explaining the test, and reviewing five problem-solving dimensions: positive problem orientation, negative problem orientation, rational style, impulsive style, and avoidant style.

### C.1.2 With Context

**Instruction.** You are a highly helpful assistant tasked with annotating the therapist's behavior. Your role is to categorize the therapist's behavior using a list of strategies provided below.

Your Task: Select one strategy from the PS Core category and one strategy from the Facilitators category that best aligns with the therapist's most recent

utterance. You are also provided with the last two utterances: the therapist's previous utterance and the client's previous utterance. You should use only these two previous utterances as context, but for annotation, you should focus only on the most recent therapist's utterance. If none of the strategies are applicable, return 'None' for that category.

**Definition.** The definition is the same as when context is not provided.

### C.2 Extracting Domain Specific Features

**Prompt.** You are a helpful assistant tasked with analyzing a therapist's utterances in a problem-solving therapy session. Your goal is to assess the utterances for the following aspects: 1. Autonomy-Supportive Language Output Labels: "Directive", "Non-Directive", or "N/A". Definitions: - Directive: - Focuses on the therapist's goals, plans, or strategies for the client. - Therapist explicitly tells the client what to do or prescribes a solution. - Prioritizes achieving a specific outcome, often determined by the therapist. - May reflect the therapist's beliefs, goals, or values rather than the client's. - Frequently uses closed-ended or leading questions to guide the client toward a particular answer or decision. - Provides answers, advice, or logical consequences rather than encouraging exploration. - Non-Directive: - Focuses on the client's goals, values, and intrinsic motivations. - Encourages the client to explore their thoughts and emotions freely without being led toward a specific conclusion. - Gives the client control over the direction of therapy and supports their decision-making process. - Avoiding judgment or imposition of therapist's opinions. - Uses open-ended questions to encourage self-reflection and insight. - Attends to the client's feelings and underlying experiences, fostering understanding and autonomy. - Accepts and explores resistance as part of the process, fostering collaboration to address challenges.



2. Self-Disclosure Output Labels: “Immediate”, “Nonimmediate”, or “N/A”. Definitions: Self-disclosure refers to instances when the therapist intentionally shares personal information, experiences, or feelings with the client. - Immediate Self-Disclosure: Sharing personal thoughts, feelings, or reactions related to the ongoing therapeutic interaction. - Nonimmediate Self-Disclosure: Discussing past experiences, personal anecdotes, or biographical information unrelated to the immediate session.

3. Open vs. Closed Questions Output Labels: “Open-Ended”, “Closed-Ended”, or “N/A”. Definitions: - Open-Ended Questions: Questions that encourage detailed, thoughtful, and expansive responses. These questions often begin with “how,” “what,” “why,” or “tell me about” and invite the client to explore their thoughts and feelings freely. - Closed-Ended Questions: Questions that elicit specific, often short responses such as “yes,” “no,” or concise information. These questions are structured and limit the scope of the client’s answer. N/A: Use this label if the utterance does not fit into either category.

4. Metaphors: Output Labels: “yes, <reasoning for metaphor detection>” (for use of metaphor) or “no” (for no use of metaphor). Definition: According to conceptual metaphor theory, metaphor facilitates a mapping of attributes or characteristics from source domain to target domain. For example, the word “invested” in the sentence “I have invested a lot of time in her” is a metaphorical expression. The source domain implied by this metaphor is the domain of money and the target domain implied by this metaphor is the domain of time. Based on the above information, decide whether this utterance uses metaphor in it and mention the phrase, source and domain of the metaphor.

Required Output Format (JSON):  
 { “autonomy”: “<string>”,  
 “self\_disclosure”: <integer>,

“question\_type”: “<string>”, “metaphor”: “<string>” }

## D Detailed Results of Classification

### D.1 LLM Classification Results

The detailed classification results for each strategy, based on the F1-scores of different LLMs, are presented in Table 7.

### D.2 Fine-tuned Models Classification Results

The detailed classification results for each strategy, based on the F1-scores of different fine-tuned transformer-based models, are available in Table 8.

## E Question Type vs. Autonomy

Figure 8 shows the co-occurrence matrix between question type and autonomy labels.

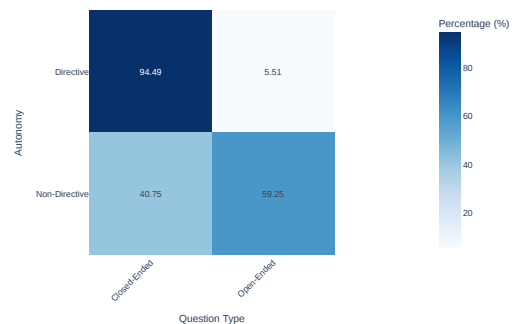


Figure 8: Co-occurrence of Question Type and Autonomy Labels

## F LIWC Results

Figures 9 and 10 present the co-occurrence of each strategy with the selected LIWC features.

	Strategies	GPT-4	GPT-4o	Llama	Yi
PS Core	Problem-solving Positive Mindset	0.83	<b>0.85</b>	0.76	0.75
	Defining Problems and Goals	0.73	<b>0.79</b>	0.72	0.63
	Generating Alternative Solutions	0.74	<b>0.79</b>	0.62	0.48
	Outcome Prediction and Planning	0.74	<b>0.84</b>	0.67	0.50
	Trying Out Solution Plan	<b>0.89</b>	0.80	0.75	0.40
Facilitators	Social Courtesies	<b>0.75</b>	0.73	0.58	0.67
	Session Management	<b>0.72</b>	0.68	0.61	0.61
	Therapeutic Engagement	0.57	<b>0.80</b>	0.69	0.64
	Test Review	0.89	<b>0.92</b>	0.79	0.70
	None	0.61	<b>0.62</b>	0.28	0.27
	weighted avg	0.68	<b>0.76</b>	0.61	0.56

Table 7: F1-scores of different LLMs for each strategy (without context)

	Strategies	DeBERTa	MentalBERT	ModernBERT	FlanT5
PS Core	Problem-solving Positive Mindset	0.81	0.78	<b>0.85</b>	0.76
	Defining Problems and Goals	0.80	0.79	<b>0.86</b>	0.79
	Generating Alternative Solutions	0.73	0.70	0.74	<b>0.75</b>
	Outcome Prediction and Planning	0.73	<b>0.77</b>	0.71	0.70
	Trying Out Solution Plan	0.80	<b>0.89</b>	0.67	0.40
Facilitators	Social Courtesies	<b>0.73</b>	0.67	0.68	0.68
	Session Management	0.66	0.63	<b>0.67</b>	0.62
	Therapeutic Engagement	<b>0.70</b>	0.68	<b>0.70</b>	0.69
	Test Review	<b>0.83</b>	0.79	0.79	0.80
	None	0.41	0.53	<b>0.58</b>	0.50
	weighted avg	0.68	0.68	<b>0.71</b>	0.67

Table 8: F1-scores of different fine-tuned models for each strategy

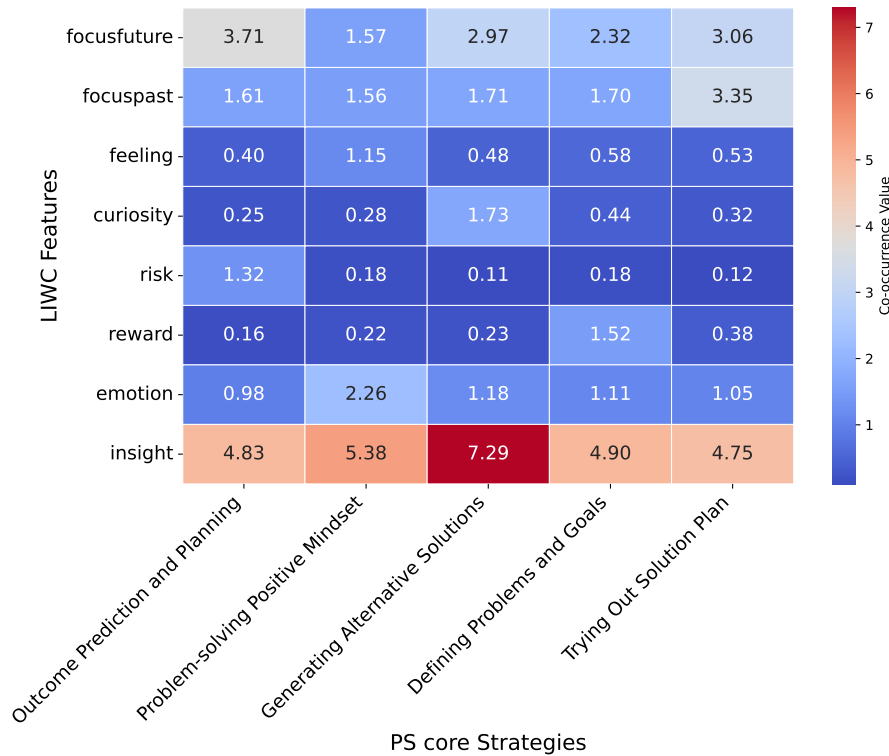


Figure 9: Co-Occurrence of PS core Strategies with Selected LIWC Features




Figure 10: Co-Occurrence of Facilitative Strategies with Selected LIWC Features